

Residential Demand Response Applications Using Batch Reinforcement Learning

F. Ruelens, B.J. Claessens, S. Vandael, B. De Schutter, R. Babuška and R. Belmans.

Abstract—Driven by recent advances in batch Reinforcement Learning (RL), this paper contributes to the application of batch RL to demand response. In contrast to conventional model-based approaches, batch RL techniques do not require a system identification step, which makes them more suitable for a large-scale implementation. This paper extends fitted Q-iteration, a standard batch RL technique, to the situation where a forecast of the exogenous data is provided. In general, batch RL techniques do not rely on expert knowledge on the system dynamics or the solution. However, if some expert knowledge is provided, it can be incorporated by using our novel policy adjustment method. Finally, we tackle the challenge of finding an open-loop schedule required to participate in the day-ahead market. We propose a model-free Monte-Carlo estimator method that uses a metric to construct artificial trajectories and we illustrate this method by finding the day-ahead schedule of a heat-pump thermostat. Our experiments show that batch RL techniques provide a valuable alternative to model-based controllers and that they can be used to construct both closed-loop and open-loop policies.

Index Terms—Batch reinforcement learning, Demand response, Electric water heater, Fitted Q-iteration, Heat pump.

I. INTRODUCTION

THE increasing share of renewable energy sources introduces the need for flexibility on the demand side of the electricity system [1]. A prominent example of loads that offer flexibility at the residential level are thermostatically controlled loads, such as heat pumps, air conditioning units, and electric water heaters. These loads represent about 20% of the total electricity consumption at the residential level in the United States [2]. In addition, their market share is expected to increase as a result of the electrification of heating and cooling [2], making them an interesting domain for optimization methods [1], [3]–[5]. Demand response programs offer demand flexibility by motivating end users to adapt their consumption profile in response to changes in the electricity price or other grid signals. The forecast uncertainty of renewable energy sources [6], combined with their limited controllability, have made demand response the topic of an extensive number of research projects [1], [7], [8] and scientific papers [3], [5], [9]–[11]. The traditional control paradigm defines the demand response problem as a model-based control problem [3], [7], [9], requiring a model of the demand response application, an optimizer, and a forecasting technique. A critical step in setting up a model-based controller comprises selecting accurate models and estimating the model parameters. This step becomes more challenging considering the heterogeneity of

the end users and their different patterns of behavior [12]. As a result, different end users are expected to have different model parameters and even different models. As such, a large-scale implementation of model-based controllers requires a stable and robust approach that is able to identify the appropriate model and the corresponding model parameters.

Reinforcement Learning (RL) [13], [14], on the other hand, is a model-free technique that requires no system identification step and no a priori knowledge. Recent developments in the field of reinforcement learning show that RL techniques can replace or supplement model-based techniques [15]. A number of recent papers provide examples of how a popular RL method, Q-learning [13], can be used for demand response [4], [10], [16], [17]. For example in [10], O'Neill *et al.* propose an automated energy management system based on Q-learning that learns how to make optimal decisions for the consumers. In [16], Henze *et al.* investigate the potential of Q-learning for the operation of commercial cold stores and in [4], Kara *et al.* use Q-learning to control a cluster of thermostatically controlled loads. In [17], Lee *et al.* propose a bias-corrected form of Q-learning to operate battery charging in the presence of volatile prices. While being a popular method, one of the fundamental drawbacks of Q-learning is its inefficient use of experiences, given that Q-learning discards the current data sample after every update. As a result, more observations are needed to propagate already known information through the state space. In order to overcome this drawback, batch RL techniques [18]–[20] can be used. In batch RL a controller estimates a control policy based on a batch of experiences. These experiences can be a fixed set [20] or can be gathered online by interacting with the environment [21]. Given that batch RL algorithms can reuse past experiences, they converge faster compared to techniques like Q-learning or SARSA. This makes batch RL techniques suitable for practical implementations, such as demand response. For example, the authors of [22] combine Q-learning with eligibility traces in order to learn the consumer and time preferences of demand response applications. In [5], the authors use a batch RL technique to schedule a cluster of electric water heaters and in [23], Vandael *et al.* use a batch RL technique to find a day-ahead consumption plan of a cluster of electric vehicles. An excellent overview of batch RL methods can be found in [21] and [24].

Inspired by the recent developments in batch RL, in particular fitted Q-iteration by Ernst *et al.* [15], this paper builds upon the existing batch RL literature and contributes to the application of batch RL techniques to residential demand response. The contributions of our paper can be summarized as follows: (1) we demonstrate how fitted Q-iteration can be extended to the situation when a forecast of the exogenous data is provided; (2) we propose a policy adjustment method

F. Ruelens, S. Vandael and R. Belmans are with the department of Electrical Engineering, KU Leuven/EnergyVille, Leuven, Belgium.

B.J. Claessens is with the energy department of VITO, Belgium.

B. De Schutter and R. Babuška are with Delft University of Technology, The Netherlands.




Fig. 1. Building blocks of a batch Reinforcement Learning (RL) controller in a demand response application.

that exploits general expert knowledge about monotonicity conditions of the control policy; (3) we introduce a model-free Monte Carlo estimator method to find a day-ahead consumption plan by making use of a novel metric based on Q-values.

This paper is structured as follows: Section II defines the building blocks of our batch RL controller. Section III formulates the problem as a Markov decision process. Section IV describes our model-free batch RL techniques for demand response. Section V demonstrates the presented techniques in a realistic demand response setting. To conclude, Section VI summarizes the results and discusses further research.

II. BUILDING BLOCKS: MODEL-FREE APPROACH

Fig. 1 presents a general overview of the different building blocks of our batch RL approach applied to a demand response setting. This paper focuses on two types of thermostatically controlled loads. The first type is a residential electric water heater with a stochastic hot-water demand [25]. The dynamic behavior of the electric water heater, used in this paper, is modeled by making use of a nonlinear stratified thermal tank model as described in [26]. Our second demand response application is a heat-pump thermostat for a residential building. The temperature dynamics of the building are modeled using a second-order equivalent thermal parameter model [27], describing the temperature dynamics of the indoor air and of the building envelope. In order to develop a practical implementation we assume that the temperature of the building envelope is a hidden state variable, and thus cannot be measured. In addition, we assume that both applications are equipped with a backup controller that guarantees the comfort and safety settings of the end users. The backup controller can be a built-in overrule mechanism that turns the application ON or OFF depending on the current state and a predefined switching logic. The operation and settings of the backup controller are assumed to be unknown, however, the batch RL controller can measure the action of the backup controller (see dashed arrow in Fig. 1).

Before the observed state information of the demand response application can be sent to the batch RL algorithm, we apply a feature extraction technique [14]. A first task of the feature extraction technique is to extract non-observable state information that is required to obtain a policy. For example, in our implementation of a heat-pump thermostat we use feature extraction to represent the temperature of the building envelope, which cannot be measured. A second task of the feature extraction technique could be to find a compact

representation of the observable state. For example, in the case of an electrical water boiler, feature extraction is used to find a compact representation of the observable state.

At the start of each day the batch RL controller constructs a control policy for the next day, given a fixed batch of transitions and cost values. The batch RL controller needs no a priori information on the model dynamics and considers its environment a black box. As a result, the batch RL controller can be applied to virtually every demand response problem or even for cluster control [5]. During the day, an exploration strategy is used online to interact with the environment and to collect new transitions that are added systematically to the batch.

The goal of this paper is to develop a model-free controller for two relevant demand response business models [1], [28]: dynamic pricing and day-ahead scheduling. The objective of the first business model is to adapt the consumption profile in response to an external price signal without violating the comfort settings of the end user. The optimal solution is a closed-loop control policy that is a function of the current and past measurements of the state. The second business model relates to the participation in the day-ahead market. The objective is to construct the day-ahead consumption plan and then follow it during the day. The goal is to minimize the cost in the day-ahead market and minimize any deviation between the day-ahead consumption plan and the actual consumption. In contrast to the solution of the first business model, the day-ahead consumption plan is a feed-forward plan for the next day, i.e. an open-loop policy, which does not depend on measurements of the state.

III. MARKOV DECISION PROCESS FORMULATION

In order to use reinforcement learning techniques, we formulate the sequential decision problem of a demand response application as a Markov decision process [14], [29]. The Markov decision process, used in this paper, is defined by its d -dimensional state space $X \subset \mathbb{R}^d$, its action space $U \subset \mathbb{R}$, its stochastic discrete-time transition function f , and its cost function ρ [18]. The optimization horizon is considered finite, comprising $T \in \mathbb{N} \setminus \{0\}$ steps, where at each discrete time step k , the state evolves as follows:

$$x_{k+1} = f(x_k, u_k, w_k) \quad \forall k \in \{1, \dots, T-1\}, \quad (1)$$

with w_k a realization of a random process drawn from a conditional probability distribution $p_{\mathcal{W}}(\cdot|x_k)$, $u_k \in U$ the control action, and $x_k \in X$ the state. Associated with each state transition, a cost c_k is given by:

$$c_k = \rho(x_k, u_k, w_k) \quad \forall k \in \{1, \dots, T\}. \quad (2)$$

The goal is to find a control policy $h^* : X \rightarrow U$ that minimizes the expected T -stage return for any state in the state space. The expected T -stage return starting from x_1 and following h^* is defined as follows:

$$J_T^{h^*}(x_1) = \mathbb{E}_{w_k \sim p_{\mathcal{W}}(\cdot|x_k)} \left[\sum_{k=1}^T \rho(x_k, h^*(x_k), w_k) \right]. \quad (3)$$

A convenient way to characterize the policy h^* is by using a state-action value function or Q-function:

$$Q^{h^*}(x, u) = \mathbb{E}_{w \sim p_W(\cdot|x)} \left[\rho(x, u, w) + J_T^{h^*}(f(x, h^*(x), w)) \right]. \quad (4)$$

The Q-function is the cumulative return starting from state x , taking action u , and following h^* thereafter. Starting from a Q-function for every state-action pair, the policy is calculated as follows:

$$h^*(x) \in \arg \min_{u \in U} Q^{h^*}(x, u), \quad (5)$$

where h^* satisfies the Bellman equation [29]. The next paragraphs give a formal description of the state space, the backup controller, and the cost function tailored to demand response.

A. State description

The state space X is spanned by a time-dependent state space component X_t , a controllable state space component X_{ph} , and an uncontrollable exogenous state space component X_{ex} [14]:

$$X = X_t \times X_{ph} \times X_{ex}. \quad (6)$$

1) *Timing*: The state space component X_t describes the part of the state space related to timing, i.e. it carries timing information that is relevant for the dynamics of the system:

$$X_t = X_t^q \times X_t^d \text{ with } X_t^q = \{1, \dots, 96\}, X_t^d = \{1, \dots, 7\}, \quad (7)$$

where $x_t^q \in X_t^q$ denotes the quarter in the day, and $x_t^d \in X_t^d$ denotes the day in the week. The rationale is that most consumer behavior tends to be repetitive and follows a diurnal pattern.

2) *Physical representation*: The controllable state space component X_{ph} represents the physical state information related to the quantities that are measured locally and that are influenced by the control actions, e.g. the indoor air temperature or the state of charge of an electric water heater:

$$x_{ph} \in X_{ph} \text{ with } \underline{x}_{ph} < x_{ph} < \overline{x}_{ph}, \quad (8)$$

where \underline{x}_{ph} and \overline{x}_{ph} denote the lower and upper bound, set to guarantee the comfort and safety of the end user.

3) *Exogenous Information*: The state description of the uncontrollable exogenous state is split into two components:

$$X_{ex} = X_{ex}^{ph} \times X_{ex}^c. \quad (9)$$

When the random disturbance w_{k+1} is independent of w_k , given x_k there is no need to include an uncontrollable exogenous state information in the state space. However, most physical processes, such as the outside temperature and solar radiation, exhibit a certain degree of autocorrelation, where the next state depends on the previous states. For this reason we include an exogenous state space component $x_{ex}^{ph} \in X_{ex}^{ph}$ in our state space description [14]. This exogenous state space component is related to the observable exogenous information that has an impact on the physical dynamics and cannot be influenced by the control actions. The second exogenous state space component $x_{ex}^c \in X_{ex}^c$ has no direct influence on the dynamics, but contains information to calculate the cost

c_k . This work assumes that a deterministic forecast of the exogenous state information related to the cost $\lambda \in \mathbb{R}^T$ is provided for the time span covering the optimization problem.

B. Backup controller

In order to develop a practical demand response technology we assume that each device is equipped with an overrule mechanism that guarantees comfort and safety constraints. The backup function $B : X \times U \rightarrow U^{ph}$ maps the requested control action $u_k \in U$ taken in state x_k to a physical control action $u_k^{ph} \in U^{ph}$:

$$u_k^{ph} = B(x_k, u_k). \quad (10)$$

The settings of the backup function B are unknown by the batch RL controller, but the resulting action u_k^{ph} can be measured (see dashed arrow in Fig. 1).

C. Cost function

In general, RL techniques do not require detailed knowledge of the cost function. However, for most demand response business models a cost function is available. This paper considers two typical cost functions related to demand response. In the dynamic pricing scenario an external price profile $\lambda \in \mathbb{R}^T$ is known deterministically at the start of the horizon. The cost function is described as:

$$c_k = u_k^{ph} \lambda_k \Delta t, \quad (11)$$

where λ_k is the electricity price at time step k and Δt is the length of a control period. The objective of the second business case is to determine a day-ahead consumption plan and to follow this plan during operation. The day-ahead consumption plan should be minimized based on day-ahead prices. In addition, any deviation between the planned consumption and actual consumption should be avoided. As such, the cost function can be written as:

$$c_k = u_k \lambda_k \Delta t + \alpha |u_k \Delta t - u_k^{ph} \Delta t|, \quad (12)$$

where u_k is the planned consumption, u_k^{ph} is the actual consumption and λ_k is the forecasted day-ahead price. The first part of (12) is the cost for buying energy at the day-ahead market. The second part defines a penalty for any deviation between the planned consumption and the actual consumption.

D. Reinforcement learning for demand response

When the description of the transition function and cost function is available, techniques that make use of the Markov decision process framework, such as approximate dynamic programming [30] or direct policy search [24], can be used to find near-optimal policies. However, in our implementation we assume that the transition function f , the backup controller B , and the underlying probability of the exogenous information are unknown. For this reason, we present a model-free batch RL approach that builds on previous theoretical work on RL, in particular fitted Q-iteration [20], expert knowledge [31], and the synthesis of artificial trajectories [18].

Algorithm 1 Fitted Q-iteration using a forecast of the exogenous data (extended FQI)

Input: $\mathcal{F} = \{(x_l, u_l, x'_l, u_l^{\text{ph}})\}_{l=1}^{\#\mathcal{F}}, \{\hat{x}_{l,\text{ex}}^{\text{ph}}\}_{l=1}^{\#\mathcal{F}}, \lambda$

```

1:  $N \leftarrow 0$ 
2: let  $\hat{Q}_0$  be zero everywhere on  $X \times U$ 
3: repeat
4:   for  $l = 1, \dots, \#\mathcal{F}$  do
5:      $c_l \leftarrow \rho(x_l, u_l^{\text{ph}}, \lambda)$ 
6:      $\hat{x}'_l \leftarrow (x_{l,t}^{\text{q}}, x_{l,t}^{\text{d}}, x_{l,\text{ph}}, \hat{x}_{l,\text{ex}}^{\text{ph}}) \triangleright$  replace the
       observed exogenous part of the next state  $x_{l,\text{ex}}^{\text{ph}}$ 
       by its forecasted value  $\hat{x}_{l,\text{ex}}^{\text{ph}}$ 
7:      $Q_{N,l} \leftarrow c_l + \min_{u \in U} \hat{Q}_{N-1}(\hat{x}'_l, u)$ 
8:   end for
9:   use regression to obtain  $\hat{Q}_N$  from
      $\mathcal{T}_{\text{reg}} = \{(x_l, u_l), Q_{N,l}\}, l = 1, \dots, \#\mathcal{F}$ 
10:  increment  $N$ 
11: until stopping criterion is reached
Output:  $Q^* = \hat{Q}_N$ 

```

IV. ALGORITHMS

Typically batch RL techniques construct policies based on a batch of tuples of the form: $\mathcal{F} = \{(x_l, u_l, x'_l, c_l)\}_{l=1}^{\#\mathcal{F}}$, where $x_l = (x_{l,t}^{\text{q}}, x_{l,t}^{\text{d}}, x_{l,\text{ph}}, x_{l,\text{ex}}^{\text{ph}})$ denotes the state at time step l and x'_l denotes the state at time step $l+1$. However, for most demand response applications, the cost function ρ is given a priori, and of the form $\rho(x_l, u_l^{\text{ph}}, \lambda)$. As such, this paper considers tuples of the form $(x_l, u_l, x'_l, u_l^{\text{ph}})$.

A. Fitted Q-iteration using a forecast of the exogenous data

Here we show how fitted Q-iteration [20] can be extended to the situation when a forecast of the exogenous state space component is provided (Algorithm 1). The algorithm iteratively builds a training set \mathcal{T}_{reg} with all state-action pairs (x, u) in \mathcal{F} as the input. The target values consist of the corresponding cost values $\rho(x, u^{\text{ph}}, \lambda)$ and the optimal Q-values, based on the approximation of the Q-function of the previous iteration, for the next states $\min_{u \in U} \hat{Q}_{N-1}(\hat{x}'_l, u)$. For a finite horizon problem the stopping criterion is reached when $N = T$, where T is the number of control periods in the optimization horizon and N denotes the iteration. It is important to note that \hat{x}'_l denotes the successor state in \mathcal{F} , where the observed exogenous state space information $x_{l,\text{ex}}^{\text{ph}}$ is replaced by its forecasted value $\hat{x}_{l,\text{ex}}^{\text{ph}}$ (line 6 in Algorithm 1). Note that in our algorithm the next state contains information on the forecasted exogenous data, whereas for standard fitted Q-iteration [20] the next state contains past observations of the exogenous data. By replacing the observed exogenous part of the next state by its forecasted value, the Q-function of the next state assumes that the exogenous information will follow its forecast. The proposed algorithm is relevant for demand response applications that are influenced by exogenous weather data. Examples of these applications are heat-pump thermostats and air conditioning units.

In principle, any regression algorithm, such as neural networks [32], can be applied in combination with fitted

Q-iteration. However, because of their robustness and fast calculation time, an extremely randomized trees ensemble method [20] is used.

B. Expert policy adjustment

Given the Q-function from Algorithm 1, a near-optimal policy can be constructed by solving (5) for every state in the state space. In this section, we show how expert knowledge on the monotonicity of the policy can be exploited to regularize the policy. The method enforces monotonicity conditions by using a convex optimization to approximate the policy, where expert knowledge is included in the form of extra constraints. These constraints can result directly from the expert or from a model-based solution. In order to define a convex optimization problem we use a fuzzy model with triangular membership functions [24] to approximate the policy. The centers of the triangular membership functions are located on an equidistant grid with N_g membership functions along each dimension of the state space. This partitioning leads to N_g^d state-dependent membership functions for each action. The parameter vector θ^* that approximates the original policy can be found by solving the following least-squares problem:

$$\theta^* \in \arg \min_{\theta} \sum_{l=1}^{\#\mathcal{F}} \left([F(\theta)](x_l) - h^*(x_l) \right)^2, \quad (13)$$

s.t. expert knowledge

where F denotes an approximation mapping of a weighted linear combination of triangular membership functions and $[F(\theta)](x)$ denotes the policy $F(\theta)$ evaluated at state x . Let h^* be the policy obtained by solving (5), given the the Q-function obtained by Algorithm 1. A more detailed description of how these triangular membership functions are defined can be found in [24]. The fuzzy approximation of the policy allows us to add expert knowledge to the policy in the form of convex constraints of the least-squares problem defined in (13), which can be solved using a convex optimization solver. Using the same notation as in [31], we can enforce monotonicity conditions along the d th dimension of state space as follows:

$$\delta_d [F(\theta)](x_d) \leq \delta_d [F(\theta)](x'_d) \quad (14)$$

for all state components $x_d \leq x'_d$ along the dimension d . If δ_d is -1 then $[F(\theta)]$ will be decreasing along the d th dimension of X , whereas if δ_d is 1 then $[F(\theta)]$ will be increasing along the d th dimension of X . Once θ^* is found, the adjusted policy \hat{h} , given this expert knowledge, can be calculated as $\hat{h}(x) = [F(\theta^*)](x)$. When the batch \mathcal{F} contains a limited number of tuples, e.g. only a few days, the expert policy adjustment method can be used to improve the quality of the policy of a demand response problem.

C. Day-ahead consumption plan

This section explains how to construct a day-ahead schedule starting from the Q-function obtained by Algorithm 1. Finding a day-ahead schedule has a direct relation to two situations: 1) a day-ahead market, where participants have to submit a

day-ahead schedule one day in advance of the actual consumption [28]; 2) a distributed optimization process, where two or more participants are coupled by a common constraint, e.g. congestion management [9]. Algorithm 2 describes a model-free Monte Carlo estimator method [18] for policy evaluation that makes use of a metric based on Q-values. The method estimates the average return of a policy by synthesizing p sequences of transitions of length T from \mathcal{F} . These p sequences can be seen as a proxy of the actual trajectories that could be obtained by simulating the policy on the given control problem. Note that since we consider a stochastic setting, p needs to be greater than 1. A sequence is grown in length by selecting a new transition among the samples of not-yet-used one-step transitions in \mathcal{F} . Each new transition is selected by minimizing a distance metric with the previously selected transition.

In [18], Fonteneau *et al.* propose the following distance metric in $X \times U : \Delta((x, x'), (u, u')) = \|x - x'\| + \|u - u'\|$, where $\|\cdot\|$ denotes the Euclidean norm. It is important to note that this metric weighs each dimensions of the state space equally. In order to overcome specifying weights to each dimension, we propose a distance metric as specified on line 8 of Algorithm 2. Here Q^* is obtained by applying Algorithm 1 and x_k^i denotes the state corresponding to the i th trajectory at time step k . The artificial trajectory P^i contains the control actions corresponding with the optimal Q-value, given the state x_k^i (see line 7). The next state $x_{l^i}^i$ is found by taking the next state of the tuple that minimizes the distance metric (see line 8). The regularization parameter ξ is a scalar that is included to penalize states that have similar Q-values, but have a large Euclidean norm in the state space. When the Q-function is strictly increasing or decreasing ξ can be set to 0. The motivation behind using Q-values instead of the Euclidean distance in $X \times U$ is that Q-values capture the dynamics of the system and, therefore, there is no need to select individual weights.

V. SIMULATIONS

This section presents the simulation results of three experiments and evaluates the performance of the proposed algorithms. We focus on two examples of flexible loads, i.e. an electric water heater and heat-pump thermostat. The first experiment evaluates the performance of extended FQI (Algorithm 1) for a heat-pump thermostat. The rationale behind using extended FQI for a heat-pump thermostat, is that the temperature dynamics of a building is influenced by exogenous weather data, which is not the case for an electric water heater. In the second experiment, we apply the policy adjustment method to an electric water heater. The final experiment uses the model-free Monte Carlo method to find a day-ahead consumption plan for a heat-pump thermostat. It should be noted that the policy adjustment method and model-free Monte Carlo method can also be applied to both heat-pump thermostat and electric water heater.

A. Thermostatically controlled loads

Here we describe the state definition and the settings of the backup controller of the electric water heater and the heat-

Algorithm 2 Model-free Monte Carlo method [18].

Input: $\mathcal{F} = \{(x_l, u_l, x'_l, u_l^{\text{ph}})\}_{l=1}^{\#\mathcal{F}}, \{\hat{x}_{l,\text{ex}}^{\text{ph}}\}_{l=1}^{\#\mathcal{F}}, \lambda, x_1, p, \xi$

- 1: $\mathcal{G} \leftarrow \mathcal{F}$
- 2: Apply Algorithm 1, to obtain Q^*
- 3: **for** $i = 1, \dots, p$ **do**
- 4: $k \leftarrow 1$
- 5: $x_k^i \leftarrow x_1$
- 6: **while** $k < T$ **do**
- 7: $u_k^i \leftarrow \arg \min_{u' \in U} Q^*(x_k^i, u')$
- 8: $\mathcal{H} \leftarrow \arg \min_{(x_l, u_l, x'_l, u_l^{\text{ph}}) \in \mathcal{G}} |Q^*(x_k^i, u_k^i) - Q^*(x_l, u_k^i)| + \xi \|x_k^i - x_l\|$
- 9: $l^i \leftarrow$ lowest index in \mathcal{G} of the transitions in \mathcal{H}
- 10: $P_k^i \leftarrow u_k^i$
- 11: $k \leftarrow k + 1$
- 12: $x_k^i \leftarrow x_{l^i}^i$
- 13: $\mathcal{G} \leftarrow \mathcal{G} \setminus \{(x_{l^i}^i, u_{l^i}^i, x'_{l^i}, u_{l^i}^{\text{ph}})\}$
- 14: **end while**
- 15: **end for**

Output: P^1, \dots, P^p

pump thermostat.

1) *Electric water heater:* We consider that the storage tank of the electric water heater is equipped with n_s temperature sensors. The full state description of the electric water heater is defined as follows:

$$x_k = (x_{k,t}^q, T_k^1, \dots, T_k^i, \dots, T_k^{n_s}), \quad (15)$$

where $x_{k,t}^q$ denotes of the current quarter in the day and T_k^i denotes the temperature measurement of the i th sensor. This work uses a feature extraction to reduce the dimensionality of the controllable state space component by replacing it with the average sensor measurement. As such the reduced state is defined as follows:

$$x_k = (x_{k,t}^q, \frac{\sum_{i=1}^{n_s} T_k^i}{n_s}). \quad (16)$$

More generic dimension reduction techniques, such as an auto-encoder network and a principle components analysis [33], [34] will be explored in future research. The logic of the backup controller of the electric water heater is defined as:

$$B(x_k, u_k) = \begin{cases} u_k^{\text{ph}} = u_{\max} & \text{if } x_{\text{soc}} \leq 30\% \\ u_k^{\text{ph}} = u_k & \text{if } 30\% < x_{\text{soc}} < 100\%. \\ u_k^{\text{ph}} = 0 & \text{if } x_{\text{soc}} \geq 100\% \end{cases} \quad (17)$$

The electric heating element of the electric water heater can be controlled with a binary control action $u_k \in \{0, u_{\max}\}$, where $u_{\max} = 2.3\text{kW}$ is the maximum power. A detailed description of the nonlinear dynamics and calculation of the state of charge x_{soc} is out of the scope of this paper and can be found in [26]. The stratified thermal tank model consists of 50 layers of which 8 are measured to construct the full state [26]. We use a set of hot-water demand profiles with a mean load of 100 l/day obtained from [25] to simulate a

realistic tap demand.

2) *Heat-pump thermostat*: Our second application considers a heat-pump thermostat that can measure the indoor air temperature, the outdoor air temperature, and the solar radiation. The full state of the heat-pump thermostat is defined as follows:

$$x_k = (x_{k,t}^q, T_{k,\text{in}}, T_{k,\text{out}}, S_k), \quad (18)$$

where the physical state space component consists of the indoor air temperature $T_{k,\text{in}}$ and the exogenous state space component contains the outside air temperature $T_{k,\text{out}}$ and the solar radiation S_k . The internal heat gains, caused by user behavior and electric appliances, cannot be measured and are not included in the state. We included a measurement of the solar radiance in the state since solar energy transmitted through windows can significantly impact the indoor temperature dynamics. In order to have a practical implementation, we consider that we cannot measure the temperature of the building envelope. Similar to [14], we used a feature extraction technique based on past state observations by including a virtual building envelope temperature, which is a running average of the past n_r air temperatures:

$$x_k = (x_{k,t}^q, T_{k,\text{in}}, \frac{\sum_{i=k-n_r}^{k-1} T_{i,\text{in}}}{n_r}, T_{k,\text{out}}, S_k). \quad (19)$$

An alternative method could be to use a model-based technique to estimate the temperature of the building envelope. The backup controller of the heat-pump thermostat is defined as follows:

$$B(x_k, u_k) = \begin{cases} u_k^{\text{ph}} = u_{\text{max}} & \text{if } T_{k,\text{in}} \leq \underline{T}_k \\ u_k^{\text{ph}} = u_k & \text{if } \underline{T}_k < T_{k,\text{in}} < \bar{T}_k \\ u_k^{\text{ph}} = 0 & \text{if } T_{k,\text{in}} \geq \bar{T}_k \end{cases}, \quad (20)$$

where \underline{T}_k and \bar{T}_k are the minimum and maximum temperature settings defined by the end user. The action space of the heat pump is discretized in 10 steps, $u_k \in \{0, \dots, u_{\text{max}}\}$, where $u_{\text{max}} = 3\text{kW}$ is the maximum power. In the simulation section we define a minimum and maximum comfort setting of 19°C and 23°C . A detailed description of the temperature dynamics of the indoor air and building envelope can be found in [27]. The exogenous information consists of the outside temperature, the solar radiation, and the internal heat gains from a location in Belgium [1], [35].

In the following experiments we define a control period of one quarter and an optimization horizon T of 96 control periods.

B. Experiment 1

The goal of the first experiment is to compare the performance of Fitted Q-Iteration (standard FQI) [20] to the performance of our extension of FQI (extended FQI, given by Algorithm 1). The objective of the considered heating system is to minimize the electricity cost of the heat pump by responding to an external price signal. The electricity prices are taken from the Belgian wholesale market [28]. We assume that a forecast of the outside temperature and solar radiance is available. Since the goal of the experiment is to assess the impact on the performance of FQI when a forecast is included,




Fig. 2. Simulation results for a heat-pump thermostat and a dynamic pricing scheme using an optimal controller, Fitted Q-Iteration (FQI) with and without forecast, and a default controller. The top plot depicts the performance metric M and the bottom plot depicts the cumulative electricity cost.

we assume perfect forecasts of the outside temperature and solar radiance. The observed state information of the controller is defined by (19) with n_r set to 3, and the cost function is given by (11). During the day, both FQI controllers use an ε -greedy exploration strategy. This exploration strategy selects a random control action with probability ε_k and follows the policy with probability $1 - \varepsilon_k$. The exploration probability ε_k is decreased on a daily basis following a harmonic sequence [30]. At the end of each day the FQI controller adds the tuples of the previous day to the batch and computes a policy for the next T time steps [36].

In order to compare the performance of the FQI controllers we define the following metric:

$$M = \frac{c_{\text{fqi}} - c_{\text{dc}}}{c_{\text{oc}} - c_{\text{dc}}}, \quad (21)$$

where c_{fqi} denotes the daily cost of the FQI controller, c_{dc} denotes the daily cost of the default controller and c_{oc} denotes the daily cost of the optimal controller. The metric M corresponds to 0 if the FQI controller obtains the same performance as the default controller and corresponds to 1 if the FQI controller obtains the same performance as the optimal controller. The default controller is a hysteresis controller that switches on when the indoor air temperature is lower than 19°C and stops heating when the indoor air temperature reaches 20°C . The optimal controller is a model-based controller that has full information on the model parameters and has perfect forecasts of all exogenous information.

The simulation results of the heating system for a simulation horizon of 80 days are depicted in Fig. 2. The top plot depicts the daily metric M and the bottom plot depicts the cumulative electricity cost of the heat-pump thermostat. The average metric M over the simulation horizon is 0.56 for standard FQI and 0.71 for extended FQI, which is an improvement of 27%. The performance gap of 0.29 between extended FQI and the optimal controller is a reasonable result given that the model dynamics and disturbances are unknown, and that exploration is included.

Extended FQI was able to decrease the total electricity cost with 19% compared to the default controller over the total simulation horizon, whereas standard FQI decreased the

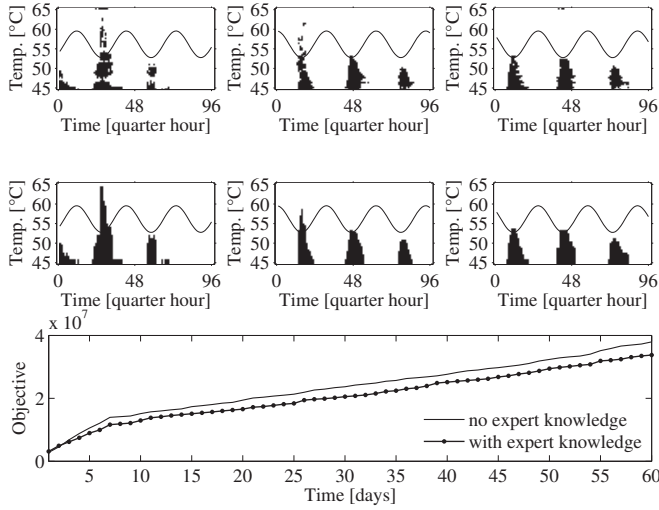


Fig. 3. Simulation results for an electric water heater and dynamic pricing scheme. The top depicts the original policies and the middle row depicts the repaired policies for day 7, 14, and 21. The price profile corresponding to each policy is depicted in the background. The bottom plot illustrates the impact of adding expert knowledge on the objective.

total electricity cost with 14%. It is important to note that the reduction of 19% is not a result of lower energy consumption, since the energy consumption increased with 4% compared to the default controller when extended FQI was used.

With these experiment we showed that we successfully extended fitted Q-iteration to incorporate a forecasts of the exogenous data.

C. Experiment 2

The following experiment demonstrates the policy adjustment method for an electric water heater. As an illustrative example we used a sinusoidal price profile. As stated in (16), the state space of an electric water heater consists of two dimensions, i.e. the time component and the average temperature. Here we enforce a monotonicity constraint along the second dimension, which contains the average temperature. The original policies after 7, 14 and 21 days obtained with fitted Q-iteration can be seen in the top row of plots of Fig. 3. It can be seen that the original policies, obtained by fitted Q-iteration, violate the monotonicity constraints along the second dimension in several states. The adjusted policies obtained by the policy adjustment method are depicted on the middle row of plots. The simulation results indicate that when the number of tuples in the batch increases, the original and adjusted policies converge. The bottom plot depicts the cumulative cost of fitted Q-iteration with and without expert knowledge over the simulation horizon. The policy adjustment method was able to reduce the total objective by 11% over 60 days. The results conclude that when the number of tuples in \mathcal{F} is small, the expert policy adjustment method can be used to improve the performance of standard fitted Q-iteration.

D. Experiment 3

The final experiment demonstrates the Model-Free Monte Carlo (MFMC) method (Algorithm 2) for finding the day-




Fig. 4. The top plot depicts the day-ahead consumption plan, actual consumption, and day-ahead price. The middle plot depicts the actual and optimal indoor temperature. The left bottom plot depicts the metric M . The daily deviation between the day-ahead consumption plan and actual consumption is given in the right bottom plot.

ahead consumption plan of a heat-pump thermostat. The state variable is defined by (19), and the cost function is defined by (12). The parameter α was set to 10^3 to penalize possible deviations between the planned consumption profile and the actual consumption. The results of the experiment are depicted in Fig. 4. The parameter p , which indicates the number of artificial trajectories, was set to 4. A day-ahead consumption plan was obtained by taking the average of these 4 trajectories. In order to assess the performance of the MFMC method, we defined an optimal controller that can exploit the model and has prescient knowledge on the internal heat gains. The top plot depicts the day-ahead consumption plan and actual consumption profile of a mature controller that contains a batch of 60 days. The corresponding indoor air temperature of the presented controller and optimal controller is depicted in the middle plot. In order to compare the daily cost of the MFMC method to the optimal controller, we define $M = c_{\text{MFMC}}/c_{\text{oc}}$, where c_{MFMC} is the daily cost of the MFMC method and c_{oc} is the daily cost of the optimal controller. The metric M for each day and the deviation between the day-ahead consumption plan and the actual followed consumption are depicted in the bottom plot. As the bottom plot indicates, the daily deviations decrease as the number of days in the batch increases. The mean metric M over the whole simulation period, including the exploration phase is 0.81. The results of the last experiment imply that the model-free Monte Carlo method can be successfully used to construct a forward consumption plan for the next day.

VI. CONCLUSION

Driven by the challenges presented by the system identification step of model-based controllers, this paper contributed to the application of model-free batch Reinforcement Learning (RL) techniques to a demand response setting. Motivated by the fact that some demand response applications, e.g a heat-pump thermostat, are influenced by exogenous weather data,

we adapted a standard batch RL technique, fitted Q-iteration, to incorporate a forecast of the exogenous data. Numerical results have been presented that indicate that the proposed extension of fitted Q-iteration was able to improve the performance of standard fitted Q-iteration by 27%. In general, batch RL techniques do not require any prior knowledge on the system behavior or the solution. However, for some demand response applications, expert knowledge about the monotonicity of the solution, i.e. the policy, can be available. As such, we presented an expert policy adjustment method that can exploit this expert knowledge. The results of an experiment with an electric water heater indicate that the policy adjustment method was able to reduce the cost objective by 11% compared to fitted Q-iteration without expert knowledge. A final challenge for model-free batch RL techniques is that of finding a day-ahead consumption plan, i.e. an open-loop solution. In order to solve this problem we presented a model-free Monte Carlo method and we successfully tested the method for finding the day-ahead consumption plan of a heat pump.

Our future research in this area will focus on employing the presented algorithms in a realistic lab environment.

ACKNOWLEDGMENTS

The authors would like to thank the reinforcement learning team from the University of Liège for their comments and suggestions.

REFERENCES

- [1] B. Dupont, P. Vingerhoets, P. Tant, K. Vanthournout, W. Cardinaels, T. De Rybel, E. Peeters, and R. Belmans, "LINEAR breakthrough project: Large-scale implementation of smart grid technologies in distribution grids," in *Proc. 3rd IEEE PES Innov. Smart Grid Technol. Conf. (ISGT Europe)*, Berlin, Germany, Oct 2012, pp. 1–8.
- [2] U.S. Energy Information Administration (EIA), "Annual energy outlook 2014 with projections to 2040," [http://www.eia.gov/forecasts/aeo/pdf/0383\(2014\).pdf](http://www.eia.gov/forecasts/aeo/pdf/0383(2014).pdf), [Online: accessed March 21, 2015].
- [3] S. Koch, J. L. Mathieu, and D. S. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in *Proc. 17th IEEE Power Sys. Comput. Conf. (PSCC)*, Stockholm, Sweden, Aug. 2011, pp. 1–7.
- [4] E. C. Kara, M. Berges, B. Krogh, and S. Kar, "Using smart devices for system-level management and control in the smart grid: A reinforcement learning framework," in *Proc. 3rd IEEE Int. Conf. on Smart Grid Commun. (SmartGridComm)*, Tainan, Taiwan, Nov. 2012, pp. 85–90.
- [5] F. Ruelens, B. Claessens, S. Vandael, S. Iacovella, P. Vingerhoets, and R. Belmans, "Demand response of a heterogeneous cluster of electric water heaters using batch reinforcement learning," in *Proc. 18th IEEE Power Sys. Comput. Conf. (PSCC)*, Wroclaw, Poland, Aug. 2014, pp. 1–8.
- [6] J. Tastu, P. Pinson, P.-J. Trombe, and H. Madsen, "Probabilistic forecasts of wind power generation accounting for geographically dispersed information," *IEEE Trans. on Smart Grid*, vol. 5, no. 1, pp. 480–489, Jan. 2014.
- [7] E. Peeters, D. Six, M. Hommelberg, R. Belhomme, and F. Bouffard, "The ADDRESS project: An architecture and markets to enable active demand," in *Proc. 6th IEEE Int. Conf. on the European Energy Market (EEM)*, Leuven, Belgium, May 2009, pp. 1–5.
- [8] F. Blik, A. van den Noort, B. Roossien, R. Kamphuis, J. de Wit, J. van der Velde, and M. Eijgelaar, "Powermatching city, a living lab smart grid demonstration," in *Proc. 1st IEEE PES Innov. Smart Grid Technol. Conf. (ISGT Europe)*, Gothenburg, Sweden, Oct. 2010, pp. 1–8.
- [9] A. Molderink, V. Bakker, M. Bosman, J. Hurink, and G. Smit, "Management and control of domestic smart grid technology," *IEEE Trans. on Smart Grid*, vol. 1, no. 2, pp. 109–119, Sept. 2010.
- [10] D. O'Neill, M. Levorato, A. Goldsmith, and U. Mitra, "Residential demand response using reinforcement learning," in *Proc. 1st IEEE Int. Conf. on Smart Grid Commun. (SmartGridComm)*, Gaithersburg, Maryland, US, Oct. 2010, pp. 409–414.
- [11] S. Vandael, B. Claessens, M. Hommelberg, T. Holvoet, and G. Deconinck, "A scalable three-step approach for demand side management of plug-in hybrid vehicles," *IEEE Trans. on Smart Grid*, vol. 4, no. 2, pp. 720–728, Nov. 2012.
- [12] J. Cigler, D. Gyalistras, J. Široký, V. Tiet, and L. Ferkl, "Beyond theory: the challenge of implementing model predictive control in buildings," in *Proc. of 11th REHVA World Congress (CLIMA)*, Czech Republic, Prague, 2013.
- [13] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, US: Cambridge Univ Press, 1998.
- [14] D. Bertsekas and J. Tsitsiklis, *Neuro-Dynamic Programming*. Nashua, NH, US: Athena Scientific, 1996.
- [15] D. Ernst, M. Glavic, F. Capitanescu, and L. Wehenkel, "Reinforcement learning versus model predictive control: a comparison on a power system problem," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 39, no. 2, pp. 517–529, 2009.
- [16] G. P. Henze and J. Schoenmann, "Evaluation of reinforcement learning control for thermal energy storage systems," *HVAC&R Research*, vol. 9, no. 3, pp. 259–275, 2003.
- [17] D. Lee and W. B. Powell, "An intelligent battery controller using bias-corrected Q-learning," in *Association for the Advancement of Artificial Intelligence*, J. Hoffmann and B. Selman, Eds. AAAI Press.
- [18] R. Fonteneau, S. A. Murphy, L. Wehenkel, and D. Ernst, "Batch mode reinforcement learning based on the synthesis of artificial trajectories," *Annals of Operations Research*, vol. 208, no. 1, pp. 383–416, 2013.
- [19] S. Adam, L. Busoniu, and R. Babuska, "Experience replay for real-time reinforcement learning control," *IEEE Trans. on Syst., Man, and Cybern., Part C: Applications and Reviews*, vol. 42, no. 2, pp. 201–212, 2012.
- [20] D. Ernst, P. Geurts, and L. Wehenkel, "Tree-based batch mode reinforcement learning," in *Journal of Machine Learning Research*, 2005, pp. 503–556.
- [21] S. Lange, T. Gabel, and M. Riedmiller, "Batch reinforcement learning," in *Reinforcement Learning: State-of-the-Art*, M. Wiering and M. van Otterlo, Eds. New York, NY, US: Springer, 2012, pp. 45–73.
- [22] Z. Wen, D. O'Neill, and H. Maei, "Optimal demand response using device-based reinforcement learning," 2015. [Online]. Available: <http://web.stanford.edu/class/ee292k/reports/ZhengWen.pdf>
- [23] S. Vandael, B. Claessens, D. Ernst, T. Holvoet, and G. Deconinck, "Reinforcement learning of heuristic EV fleet charging in a day-ahead electricity market," *IEEE Trans. on Smart Grid*, March 2015.
- [24] L. Busoniu, R. Babuska, B. De Schutter, and D. Ernst, *Reinforcement Learning and Dynamic Programming Using Function Approximators*. CRC Press, 2010.
- [25] U. Jordan and K. Vajen, "Realistic domestic hot-water profiles in different time scales," *Report for IEA-SHC Task*, vol. 26, 2001.
- [26] K. Vanthournout, R. D'hulst, D. Geysen, and G. Jacobs, "A smart domestic hot water buffer," *IEEE Trans. on Smart Grid*, vol. 3, no. 4, pp. 2121–2127, Dec. 2012.
- [27] D. Chassin, K. Schneider, and C. Gerkensmeyer, "GridLAB-D: An open-source power systems modeling and simulation environment," in *Proc. IEEE Transmission and Distribution Conf. and Expos.*, Chicago, IL, US, April 2008, pp. 1–5.
- [28] "Belpex - belgian power exchange," <http://www.belpex.be/>, [Online: accessed March 21, 2015].
- [29] R. Bellman, *Dynamic Programming*. New York, NY, US: Dover Publications, Inc., 2003.
- [30] W. Powell, *Approximate Dynamic Programming: Solving The Curses of Dimensionality*, 2nd edition. Hoboken, NJ, US: Wiley-Blackwell, 2011.
- [31] L. Busoniu, D. Ernst, R. Babuska, and B. De Schutter, "Exploiting policy knowledge in online least-squares policy iteration: An empirical study," *Automation, Computers, Applied Mathematics*, vol. 19, no. 4, 2010.
- [32] M. Riedmiller, "Neural fitted Q-iteration—first experiences with a data efficient neural reinforcement learning method," in *Machine Learning: ECML 2005*. Springer, 2005, pp. 317–328.
- [33] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [34] S. Lange and M. Riedmiller, "Deep auto-encoder neural networks in reinforcement learning," in *Proc. IEEE 2010 Int. Joint Conf. on Neural Networks (IJCNN)*, Barcelona, Spain, July 2010, pp. 1–8.
- [35] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, W. F. Buhl, Y. J. Huang, C. O. Pedersen, R. K. Strand, R. J. Liesen, D. E. Fisher, M. J. Witte et al., "EnergyPlus: creating a new-generation building energy simulation program," *Energy and Buildings*, vol. 33, no. 4, pp. 319–331, 2001.

- [36] T. Gabel, C. Lutz, and M. Riedmiller, "Improved neural fitted Q-iteration applied to a novel computer gaming and learning benchmark," in *IEEE Symp. on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, Paris, France, April 2011, pp. 279–286.